Non-Greedy L21-Norm Maximization for Principal Component Analysis

Feiping Nie, Heng Huang

arXiv:1603.08293v1 [cs.LG] 28 Mar 2016

***Abstract*— Principal Component Analysis (PCA) is one of the most important unsupervised methods to handle high-dimensional data. However, due to the high computational complexity of its eigen decomposition solution, it hard to apply PCA to the large-scale data with high dimensionality. Meanwhile, the squared L2-norm based objective makes it sensitive to data outliers. In recent research, the L1-norm maximization based PCA method was proposed for efficient computation and being robust to outliers. However, this work used a greedy strategy to solve the eigen vectors. Moreover, the L1-norm maximization based objective may not be the correct robust PCA formulation, because it loses the theoretical connection to the minimization of data reconstruction error, which is one of the most important intuitions and goals of PCA. In this paper, we propose to maximize the L21-norm based robust PCA objective, which is theoretically connected to the minimization of reconstruction error. More importantly, we propose the efficient non-greedy optimization algorithms to solve our objective and the more general L21-norm maximization problem with theoretically guaranteed convergence. Experimental results on real world data sets show the effectiveness of the proposed method for principal component analysis.**

***Index Terms*— Principal component analysis, robust dimensionality reduction, L21-norm maximization.**

I. Introduction

In many real-world applications, the dimensionality of data are very high. Directly handle the high-dimensional data is computationally expensive. At the same time, the performance could be poor because the number of available data is often limited and the noise in the data would increase dramatically when the dimensionality increases. To solve these problems, dimensionality reduction is one of the most important and effective methods. Among the dimensionality reduction algorithms, Principal Component Analysis (PCA) [1] is one of the most widely used algorithms due to its simplicity and effectiveness. The main goal of PCA is to preserve the structure of the original data in the projected low-dimensional space. To this end, given a data set, PCA finds a projection matrix to minimize the reconstruction error of the projected data points under this projection matrix.

The traditional PCA has been successfully applied in many problems [2] in the past decades. However, the traditional PCA algorithm has several drawbacks. First, it need perform Singular Vector Decomposition (SVD) on input data matrix or eigen-decomposition on the covariance matrix, which is computationally expensive and difficult to be used when the number and the dimensionality of data are both very high. Second, it is sensitive to data outliers, because its objective function is intrinsically based on squared L2-norm and the outliers with large variation values can be exaggerated by the squared L2-norm. Many recent research works [3], [4], [5], [6], [7], [8] have devoted effort to alleviate this problem and improve the robustness to outliers. [3], [6] proposed to find the subspace such that the sum of L1-norm distances of data points to the subspace is minimized. Although the robustness to outliers is improved in these methods, their algorithms are computationally expensive. Moreover, the used L1-norm in objective is not invariant to rotation. Thus, the performance is usually poor when their L1-norm based PCA is combined with K-means clustering [7]. To solve this problem, the R1-PCA was proposed with rotational invariant property and demonstrated good performance [7]. However, the R1-PCA iteratively performs the subspace iteration algorithm [9] in the high dimensional original space, which is computationally expensive. The extension of R1-PCA to tensor version can be found in [10].

Recently, a PCA method based on L1-norm maximization was proposed in [11], and a similar work can be found in [12]. This method is invariant to rotation and is also robust to outliers. An efficient algorithm was proposed to solve the L1-norm maximization problem in [11]. This algorithm only need perform matrix-vector multiplication, and thus can be applied in the case that both the number and the dimensionality of data are very high. Several works on its tensor version and supervised version can be found in [13], [14], [15]. Due to the difficulty of directly solving the L1-norm maximization problem, all these works use a greedy strategy to solve it. Specifically, the projection directions are optimized one by one sequentially. Such a kind of greedy method is easy to get stuck in a local solution.

Moreover, the L1-norm maximization based PCA method is not theoretically connected to minimization of the reconstruction error, which is the important goal of traditional PCA. In this paper, we propose a novel principal component analysis method based on the L21-norm maximization. The proposed method is robust to data outliers and also invariant to rotation. More importantly, our new method is theoretically connected to the minimization of reconstruction error, and thus is more suitable for principal component analysis than previous method in [11]. To solve the derived L21-norm PCA objective, we propose a new *non-greedy* and efficient optimization algorithm to optimize all the projection directions simultaneously. Meanwhile, our algorithm will be extended to solve the more general maximization problems. We provide the theoretical analysis to guarantee the convergence of our algorithms. All experimental results on real world data sets show that the proposed method is effective for principal component analysis, and always obtains smaller reconstruction error than the method in [11] under the same reduced dimension.

The rest of this paper is organized as follows: We give a brief review of the related work in Section 2. In Section 3, we propose the L21-norm maximization based principal component analysis and solve the derived optimization problem through a new non-greedy and efficient algorithm which can solve the more general L21-norm maximization problem. In Section 4, we extend our

Feiping Nie and Heng Huang are with the Department of Computer Science and Engineering, University of Texas at Arlington, USA. Email: feipingnie@gmail.com, heng@uta.edu

algorithm to solve the general maximization problem which can be used to derive solutions for many other statistical learning models. In Section 5, we present experimental results to verify the effectiveness of the proposed method. Finally, we draw the conclusions in Section 6.

II. Related work

Given data $X = [x_1, x_2, \cdots, x_n] \in \Re^{d\times n}$, where d and n are the dimensionality and number of data points respectively, without loss of generality, we can assume the data $\{x_i\}_{i=1}^n$ are centralized, *i.e.*, $\sum_{i=1}^n x_i = 0$.

We denote the projection matrix $W = [w_1, w_2, \cdots, w_m] \in \Re^{d\times m}$. The traditional PCA method minimizes the reconstruction error under the projected subspace, which is to solve the following optimization problem:

$$\min_{W^TW=I} \sum_{i=1}^{n} \left\| x_i - WW^Tx_i \right\|_2^2, \tag{1}$$

where I is the identity matrix, $\|\cdot\|_2$ is the L2-norm of vector. Equivalently, the traditional PCA method can also be formulated as maximizing the variance of data in the projected subspace, which is to solve the following optimization problem:

$$\max_{W^TW=I} Tr(W^TS_tW) = \max_{W^TW=I} \sum_{i=1}^{n} \left\| W^Tx_i \right\|_2^2, \tag{2}$$

where $S_t = XX^T$ is the covariance matrix and $Tr(\cdot)$ is the trace operator of a matrix. The equivalence of Eq. (1) and Eq. (2) is based on the following equation for any matrix W with $W^TW = I$:

$$\sum_{i=1}^{n} \left\| x_i - WW^Tx_i \right\|_2^2 + \sum_{i=1}^{n} \left\| W^Tx_i \right\|_2^2 = \sum_{i=1}^{n} \|x_i\|_2^2 \tag{3}$$

Based on Eq. (1), R1-PCA was proposed to solve the following problem [7]:

$$\min_{W^TW=I} \sum_{i=1}^{n} \left\| x_i - WW^Tx_i \right\|_2. \tag{4}$$

R1-PCA minimizes the L2-norm loss instead of the squared L2-norm loss in traditional PCA, and thus the robustness to outliers is improved. The other important property of R1-PCA is that it is invariant to rotation, which is also preserved by traditional PCA.

Motivated by Eq. (2), a recent work named PCA-L1 [11] was proposed to maximize the L1-norm instead of the squared L2-norm in traditional PCA by solving the following problem:

$$\max_{W^TW=I} \sum_{i=1}^{n} \left\| W^Tx_i \right\|_1, \tag{5}$$

where $\|\cdot\|_1$ is the L1-norm of vector. PCA-L1 also has the rotation invariant property. Directly solving this problem is difficult, thus the author used a greedy strategy to solve it. Specifically, the m projection directions $\{w_1, w_2, \cdots, w_m\}$ are optimized one by one. The first projection direction w_1 is optimized by solving:

$$\max_{w_1^Tw_1=1} \sum_{i=1}^{n} |w_1^Tx_i|. \tag{6}$$

After the $(k-1)$-th projection direction w_{k-1} has been obtained, the data matrix X is transformed to $X = X - w_{k-1}(w_{k-1})^TX$, and then the k-th projection direction w_k is optimized by solving:

$$\max_{w_k^Tw_k=1} \sum_{i=1}^{n} |w_k^Tx_i|. \tag{7}$$

In this greedy method, the problem (7) is the key function to be solved for each k. The work in [11] proposed an iterative algorithm to solve this problem. In order to guarantee the algorithm converges to a local maximum, the algorithm adds an additional judgement after the convergence to w_k^t. If there exists i such that $(w_k^t)^Tx_i = 0$, then let $w_k^t = (w_k^t + \triangle w)/\|w_k^t + \triangle w\|_2$ and re-run the iterative algorithm, where $\triangle w$ is a small nonzero random vector. However, such operation might make the algorithm interminable. For example, if there is a data point x that exactly locates on the mean of the data set, then x will be zero after centralization. As a result, $(w^t)^Tx$ is always zero for any w^t). Moreover, it is possible that there exists i such that $(w_k^t)^Tx_i = 0$ at the global maximum. In this case, the algorithm doesn't have the chance to find the global maximum.

III. Principal component analysis with non-greedy L21-Norm maximization

A. L21-norm principal component analysis

Motivated by Eq. (2), we propose to solve the following problem:

$$\max_{W^TW=I} \sum_{i=1}^{n} \|W^Tx_i\|_2 = \max_{W^TW=I} \|X^TW\|_{2,1}, \tag{8}$$

where $\|\cdot\|_{2,1}$ is the L21-norm of a matrix defined as $\|M\|_{2,1} = \sum_i (\sum_j m_{ij}^2)^{\frac{1}{2}}$. Contrast to the name of PCA-L1 in [11] that solves problem (5), we call our PCA method with solving problem (8) as PCA-L21. The PCA-L21 maximizes the L2-norm instead of the squared L2-norm in PCA, and thus the robustness to outliers is also improved. Similarly to R1-PCA and PCA-L1, PCA-L21 is also a rotation invariant method.

It is conjectured in [11] that problem (4) and problem (5) are closely related. However, no theoretical analysis was provided in [11] and it seems not the case according to our extensively experimental results. In contrast, we will show from both theoretical and experimental results that the proposed problem (8) is indeed closely related to the problem (4), thus PCA-L21 is more suitable for the principal component analysis than PCA-L1.

First, we have the following lemma:

Lemma 1: If $a^2 + b^2 = c^2$, then $|c| \le |a| + |b| \le \sqrt{2}\,|c|$.

Proof: Starting from the condition, we have

$$\begin{aligned} & a^2 + b^2 = c^2 \\ \Rightarrow\quad & a^2 + b^2 + 2\,|a|\,|b| \ge c^2 \\ \Rightarrow\quad & (|a| + |b|)^2 \ge c^2 \\ \Rightarrow\quad & |a| + |b| \ge |c|\,. \end{aligned} \tag{9}$$

On the other hand, we have

$$\begin{aligned} & (|a| - |b|)^2 \ge 0 \\ \Rightarrow\quad & (|a| + |b|)^2 \le 2(a^2 + b^2) \\ \Rightarrow\quad & (|a| + |b|)^2 \le 2c^2 \\ \Rightarrow\quad & |a| + |b| \le \sqrt{2}\,|c| \end{aligned} \tag{10}$$

Combining Eq. (9) and Eq. (10), we complete the proof.

□

According to Eq. (3) and Lemma 1, we have the following relationship:

$$\sum_{i=1}^{n}\|x_i\|_2 \leq \sum_{i=1}^{n}\left\|x_i - WW^T x_i\right\|_2 + \sum_{i=1}^{n}\left\|W^T x_i\right\|_2 \leq \sqrt{2}\sum_{i=1}^{n}\|x_i\|_2, \tag{11}$$

which can be written in matrix form as

$$\|X\|_{2,1} \leq \|X^T - X^T WW^T\|_{2,1} + \|W^T X\|_{2,1} \leq \sqrt{2}\,\|X\|_{2,1}\,. \tag{12}$$

Therefore, the proposed problem (8) is theoretically connected to the problem (4), which indicates that maximizing the L21 norm as in problem (8) also makes sense to minimize the reconstruction error, and thus is suitable for the principal component analysis.

To solve the problem (8), we first propose an efficient algorithm to solve the more general L21-norm maximization problem. Utilizing this general algorithm, we can solve the problem (8) directly without using the greedy strategy as in [11].

B. Efficient algorithm to solve the general L21-norm maximization problem

Consider a general L21-norm maximization problem as follows:

$$\max_{v\in\mathcal{C}} f(v) + \sum_i \|g_i(v)\|_2, \tag{13}$$

where $f(v)$ is an arbitrary scatter-output function, $g_i(v)$ (for each i) is an arbitrary vector-output function, and $v \in \mathcal{C}$ is an arbitrary constraint. We assume that the objective in problem (13) has an upper bound.

We re-write the problem (13) as the following problem:

$$\max_{v\in\mathcal{C}} f(v) + \sum_i (\alpha_i)^T g_i(v), \tag{14}$$

where

$$\alpha_i = \begin{cases} \frac{g_i(v)}{\|g_i(v)\|_2} & if\ \ \|g_i(v)\|_2 \neq 0\ ; \\ \mathbf{0} & if\ \ \|g_i(v)\|_2 = 0\ . \end{cases} \tag{15}$$

Note that α_i depends on v and thus is also an unknown variable. Based on Eqs. (14) and (15), we propose an iterative algorithm to solve the problem (13). The algorithm is described in Algorithm 1. In each iteration, α_i is updated by current solution v, and the solution v is updated with the updated α_i. The iterative procedure is repeated till the algorithm converges.

Initialize $v^1 \in \mathcal{C}$, $t = 1$;
while *not converge* **do**
1. For each i, calculate α_i^t according to Eq. (15) ;
2. $v^{t+1} = \arg\max_{v\in\mathcal{C}} f(v) + \sum_i (\alpha_i^t)^T g_i(v)$;
3. $t = t + 1$;
end
Output: v^t.

Algorithm 1: An efficient algorithm to solve the general L21-norm maximization problem (13).

Next, we prove that the proposed iterative algorithm will monotonically increase the objective function value of the problem (13) in each iteration, and will converge to a local solution.

The convergence of the Algorithm 1 is demonstrated in the following theorem:

Theorem 1: The Algorithm 1 monotonically increases the objective function value of the problem (13) in each iteration.

Proof: For each iteration t, according to the Step 2 in Algorithm 1, we have

$$f(v^{t+1}) + \sum_i (\alpha_i^t)^T g_i(v^{t+1}) \geq f(v^t) + \sum_i (\alpha_i^t)^T g_i(v^t). \tag{16}$$

On the other hand, for each i, according to the Cauchy-Schwarz inequality, we know:

$$\|g_i(v^t)\|_2 \|g_i(v^{t+1})\|_2 \geq g_i(v^t)^T g_i(v^{t+1}).$$

Based on this inequality, Eq. (15), and

$$\|g_i(v^t)\|_2 - (\alpha_i^t)^T g_i(v^t) = 0,$$

we have

$$\begin{aligned} & \|g_i(v^t)\|_2 \|g_i(v^{t+1})\|_2 \geq g_i(v^t)^T g_i(v^{t+1}) \\ \Rightarrow\ & \|g_i(v^{t+1})\|_2 \geq (\alpha_i^t)^T g_i(v^{t+1}) \\ \Rightarrow\ & \|g_i(v^{t+1})\|_2 - (\alpha_i^t)^T g_i(v^{t+1}) \geq 0 \\ \Rightarrow\ & \|g_i(v^{t+1})\|_2 - (\alpha_i^t)^T g_i(v^{t+1}) \geq \|g_i(v^t)\|_2 - (\alpha_i^t)^T g_i(v^t). \end{aligned}$$

The above inequality holds for every i, thus we have

$$\begin{aligned} & \sum_i \|g_i(v^{t+1})\|_2 - \sum_i (\alpha_i^t)^T g_i(v^{t+1}) \\ & \geq \sum_i \|g_i(v^t)\|_2 - \sum_i (\alpha_i^t)^T g_i(v^t)\,. \end{aligned} \tag{17}$$

Combining Eq. (16) and Eq. (17), we arrive at

$$f(v^{t+1}) + \sum_i \|g_i(v^{t+1})\|_2 \geq f(v^t) + \sum_i \|g_i(v^t)\|_2\,. \tag{18}$$

Thus, the Algorithm 1 monotonically increases the objective of the problem (13) in each iteration t. □

As the objective of the problem (13) has an upper bound, Theorem 1 indicates that the Algorithm 1 converges. The following theorem shows that the Algorithm 1 will converge to a local solution.

Theorem 2: The Algorithm 1 will converge to a local solution of the problem (13).

Proof: The Lagrangian function of the problem (13) is

$$\mathcal{L}(v,\lambda) = f(v) + \sum_i \|g_i(v)\|_2 - r(v,\lambda), \tag{19}$$

where $r(\lambda, v)$ is the Lagrangian term to encode the constraint $v \in \mathcal{C}$ in problem (13).

Taking the derivative[1] of $\mathcal{L}(v,\lambda)$ *w.r.t.* v, and setting the derivative to zero, we have:

$$\frac{\partial \mathcal{L}(v,\lambda)}{\partial v} = \frac{\partial f(v)}{\partial v} + \sum_i J_i(v)\alpha_i - \frac{\partial r(v,\lambda)}{\partial v} = \mathbf{0}, \tag{20}$$

where α_i is defined in Eq. (15) and $J_i(v)$ is a matrix with the (j,k)-th element as $\frac{\partial g_i^k(v)}{\partial v_j}$, $g_i^k(v)$ denotes the k-th element of the vector $g_i(v)$.

Suppose the Algorithm 1 converges to a solution v^*. From the Step 2 in Algorithm 1 we have

$$v^* = \arg\max_{v\in\mathcal{C}} f(v^*) + \sum_i (\alpha_i^*)^T g_i(v^*). \tag{21}$$

[1]When $g_i(v) = \mathbf{0}$, then $\mathbf{0}$ is a subgradient of function $\|g_i(v)\|_2$, so the α_i defined in Eq. (15) is the gradient or a subgradient of the function $\|g_i(v)\|_2$ in all the cases.

According to the KKT condition [16] of the problem (21), we know that the solution v^* for problem (21) satisfies Eq. (20). Note that Eq. (20) is the KKT condition of the problem (13), hence the solution v^* satisfies the KKT condition of the problem (13). Therefore, the converged solution v^* is a local solution of the problem (13).

□

C. Non-greedy maximization algorithm to solve L21-Norm Principal component analysis

Obviously the proposed problem (8) is a special case of the problem (13), thus we can use the proposed Algorithm 1 to solve the objective of L21-Norm PCA.

In Algorithm 1, Step 2 is the key step. Thus, to solve the problem (8), the key step is to solve the following problem:

$$\max_{W^TW=I} \sum_{i=1}^{n} \alpha_i^T W^T x_i \,, \tag{22}$$

where the vector $\alpha_i \in \Re^{m\times 1}$ is defined as:

$$\alpha_i = \begin{cases} \frac{W^T x_i}{\|W^T x_i\|_2} & if\ \|W^T x_i\|_2 \neq 0\,; \\ \mathbf{0} & if\ \|W^T x_i\|_2 = 0\,. \end{cases} \tag{23}$$

Denoting the matrix $M = \sum_{i=1}^{n} x_i \alpha_i^T \in \Re^{d\times m}$, we can re-write the problem (22) as:

$$\max_{W^TW=I} Tr(W^T M). \tag{24}$$

Suppose the SVD result of M is $M = U\Lambda V^T$, then $Tr(W^T M)$ can be re-written as:

$$\begin{aligned} & Tr(W^T M) \\ = \ & Tr(W^T U \Lambda V^T) \\ = \ & Tr(\Lambda V^T W^T U) \\ = \ & Tr(\Lambda Z) \\ = \ & \sum_i \lambda_{ii} z_{ii} \,, \end{aligned} \tag{25}$$

where $Z = V^T W^T U$, λ_{ii} and z_{ii} are the (i,i)-th element of matrix λ and Z, respectively.

Note that Z is an orthonormal matrix, *i.e.* $Z^T Z = I$, so $z_{ii} \leq 1$. On the other hand, $\lambda_{ii} \geq 0$, since λ_{ii} is singular value of M. Therefore, $Tr(W^T M) = \sum_i \lambda_{ii} z_{ii} \leq \sum_i \lambda_{ii}$, and when $z_{ii} = 1$ $(1 \leq i \leq c)$, the equality holds. That is to say, $Tr(W^T M)$ reaches the maximum, when $Z = I$. Recall that $Z = V^T W^T U$, thus the optimal solution to the problem Eq. (24) is

$$W = UZ^T V^T = UV^T. \tag{26}$$

Based on the Algorithm 1, the algorithm of PCA-L21 to solve problem (8) is described in Algorithm 2. According to Theorem 2, we can obtain a local solution with the algorithm. Contrast to the PCA-L1 algorithm in [11], the PCA-L21 algorithm directly solves the projection matrix W (*i.e.* optimizes all projection directions simultaneously), but the PCA-L1 algorithm solves the projection directions one by one using a greedy strategy.

From Algorithm 2 we can see that the computational complexity of the algorithm is $O(ndmt)$, where n, d, m, t is the number of data points, the original dimensionality, the reduced dimensionality and the iteration number, respectively. In practice, the algorithm usually converges in ten iterations. Therefore, the computational complexity of the algorithm is linear *w.r.t.* either data number or data dimension, which indicates the algorithm is applicable in the case that both data number and data dimension are high. If the data are sparse, the computational complexity is further reduced to $O(nsmt)$, where s is the averaged number of nonzero elements in one data point.

Input: X, m, where X is centralized
Initialize $W^1 \in \Re^{d\times m}$ such that $W^T W = I$, $t = 1$;
while *not converge* **do**
  1. For each i, calculate α_i^t according to Eq. (23), $M = \sum_{i=1}^{n} x_i \alpha_i^T$;
  2. Calculate the SVD of M as $M = U\Lambda V^T$, Let $W^{t+1} = UV^T$;
  3. $t = t + 1$;
end
Output: $W^t \in \Re^{d\times m}$.

Algorithm 2: The non-greedy optimization algorithm to solve the L21-norm principal component analysis.

D. Kernel and tensor extensions of L21-norm PCA

Similar to traditional PCA, the proposed PCA-L21 is also a linear method, and is not suitable to handle data under the non-Gaussian distribution. A popular technique to deal with this problem is to extend the linear method to kernel method. Obviously, the PCA-L21 is invariant to rotation and shift, so this linear method satisfies the conditions of the general kernelization framework in [17]. Thus, the PCA-L21 can be kernelized using this framework. Specifically, the given data are transformed by KPCA [18], and then Algorithm 2 is performed on the transformed data.

Another problem of PCA is that it can only handle data points with vector format. For high-order tensor data, we have to vectorize the data to very high-dimensional vectors before applying PCA. This approach destroys the spacial information of tensor data and makes the computational burden very heavy. A popular technique to deal with this problem is to extend the vector based method to tensor based method. As the problem (8) in PCA-L21 only includes linear operator $W^T x_i$, it can be easily extended to the tensor method to handle high-order data directly. For simplicity, we only briefly discuss the case of 2D tensor, the higher order tensor cases can be readily extended by replacing the linear operator $W^T x_i$ with tensor operator [19].

Given data $X = [X_1, X_2, \cdots, X_n] \in \Re^{r\times c\times n}$, where each data $X_i \in \Re^{r\times c}$ is a 2D matrix, n is the number of data points, we assume that $\{X_i\}_{i=1}^n$ are centralized, *i.e.*, $\sum_{i=1}^n X_i = \mathbf{0}$. To handle the tensor case, linear operator $W^T x_i$ is replaced by $U^T X_i V$, where $U \in \Re^{r\times k_1}$ and $V \in \Re^{c\times k_2}$ are two projection matrices ($k_1 < r$ and $k_2 < c$ are the reduced dimensions of two projection subspaces). Correspondingly, the problem (8) becomes:

$$\max_{U^TU=I_{k_1}, V^TV=I_{k_2}} \sum_{i=1}^{n} \|U^T X_i V\|_F, \tag{27}$$

where $\|\cdot\|_F$ is the Frobenius norm of matrix. Similar to other tensor methods, problem (27) can be solved by alternative optimization technique. Specifically, when fix U, the problem

(27) is reduced to the problem (8), and thus the V can be optimized by Algorithm 2. In turn, U can also be optimized by Algorithm 2 when fix V. The procedure is iteratively performed until converges.

IV. The Extension of our algorithm for general maximization problem

Besides solving the L21-norm maximization problem in above section, to provide the useful and efficient algorithm for related research problems, we extend our idea to solve the more general maximization problem as follows:

$$\max_{v\in\mathcal{C}} f(v) + \sum_i h_i(g_i(v)), \tag{28}$$

where $f(v)$ is an arbitrary scatter-output function, $g_i(v)$ (for each i) is an arbitrary scatter, vector, or matrix-output function, and h_i (for each i) is an arbitrary **convex** function, $v \in \mathcal{C}$ is an arbitrary constraint. We assume that the objective in problem (28) has an upper bound.

We propose an iterative algorithm to solve the problem (13). The algorithm is described in Algorithm 3. Similar to Algorithm 1, in each iteration, α_i is updated by current solution v, and the solution v is updated with the updated α_i. The iterative procedure is repeated till the algorithm converges.

Initialize $v^1 \in \mathcal{C}$, $t = 1$;
while *not converge* **do**
1. For each i, calculate $\alpha_i^t = h_i'(g_i(v^t)) = \frac{\partial h_i(g_i(v^t))}{\partial (g_i(v^t))}$;
2. $v^{t+1} = \arg\max_{v\in\mathcal{C}} f(v) + \sum_i Tr((\alpha_i^t)^T g_i(v))$;
3. $t = t + 1$;

end
Output: v^t.

Algorithm 3: An efficient algorithm to solve the more general maximization problem (28).

The convergence of the Algorithm 3 is guaranteed by the following theorem:

Theorem 3: The Algorithm 3 monotonically increases the objective of the problem (28) in each iteration.

Proof: For each iteration t, according to the Step 2 in Algorithm 3, we have

$$f(v^{t+1}) + \sum_i Tr((\alpha_i^t)^T g_i(v^{t+1})) \geq f(v^t) + \sum_i Tr((\alpha_i^t)^T g_i(v^t)). \tag{29}$$

For each i, since h_i is convex, according to the property of convex function, we know $h_i(g_i(v^{t+1})) - h_i(g_i(v^t)) \geq Tr((g_i(v^{t+1}) - g_i(v^t))^T h_i'(g_i(v^t))$. According to the definition of α_i^t in Step 1, we have

$$\begin{aligned} &h_i(g_i(v^{t+1})) - Tr((\alpha_i^t)^T g_i(v^{t+1})) \\ &\geq h_i(g_i(v^t)) - Tr((\alpha_i^t)^T g_i(v^t)). \end{aligned} \tag{30}$$

Combining Eq. (29) and Eq. (30), we arrive at

$$f(v^{t+1}) + \sum_i h_i(g_i(v^{t+1})) \geq f(v^t) + \sum_i h_i(g_i(v^t)). \tag{31}$$

Thus the Algorithm 3 monotonically increases the objective of the problem (28) in each iteration. □

Because the objective of the problem (28) has an upper bound, Theorem 3 indicates that the Algorithm 3 converges. The following theorem shows that the Algorithm 3 will converge to a local solution of the problem (28).

Theorem 4: The Algorithm 3 will converge to a local solution of the problem (28).

Proof: The Lagrangian function of the problem (28) is

$$\mathcal{L}(v,\lambda) = f(v) + \sum_i h_i(g_i(v)) - r(v,\lambda), \tag{32}$$

where $r(\lambda, v)$ is the Lagrangian term to encode the constraint $v \in \mathcal{C}$ in problem (28).

Taking the derivative [2] of $\mathcal{L}(v,\lambda)$ *w.r.t.* v, and setting the derivative to zero, we have:

$$\frac{\partial \mathcal{L}(v,\lambda)}{\partial v} = \frac{\partial f(v)}{\partial v} + \sum_i \frac{Tr\left((\alpha_i)^T \partial g_i(v)\right)}{\partial v} - \frac{\partial r(v,\lambda)}{\partial v} = \mathbf{0}, \tag{33}$$

where $\alpha_i = h_i'(g_i(v))$.

Suppose the Algorithm 3 converges to a solution v^*. From the Step 2 in Algorithm 3 we have

$$v^* = \arg\max_{v\in\mathcal{C}} f(v^*) + \sum_i Tr((\alpha_i^*)^T g_i(v^*)). \tag{34}$$

According to the KKT condition of the problem (34), we know that the solution v^* to problem (34) satisfies Eq. (33). Note that Eq. (33) is the KKT condition of the problem (28), so the converged solution v^* of Algorithm 3 satisfies the KKT condition of the problem (28). Therefore, the converged solution v^* is a local solution of the problem (28). □

Algorithm 3 is very useful to solve the general maximization problems. For example, we can directly use the algorithm to solve the following two important problems:

$$\max_{V^TV=I} Tr(V^TAV)\,, \tag{35}$$

$$\max_{v^Tv=1, \|v\|_0\leq k} v^TAv\,. \tag{36}$$

It is interesting to point out that the derived algorithms for Eq. (35) and Eq. (36) based on Algorithm 3 are exactly the classical power method (or subspace iteration method) and the recently proposed truncated power method [20], respectively.

V. Experiments

In this section, we will present the experimental results to demonstrate the effectiveness of the proposed PCA-L21 compared to traditional PCA, R1-PCA and PCA-L1.

A. Reconstruction errors with occlusions

We use six image benchmark data sets to perform our experiments. A brief descriptions of the data sets are shown in Table I, and the samples from each data sets are shown in Figure 1.

In each image data set, we randomly select 10, 20, 30 percent images respectively, and each selected image is occluded with a randomly located square. The width of these squares are the half of the width of the images.

[2] When $x = \mathbf{0}$, then $\mathbf{0}$ is a subgradient of function $\|x\|_2$, so α_i is the gradient or a subgradient of the function $\|x\|_2$ in all the cases, where $\alpha_i = x/\|x\|_2$ if $\|x\|_2 \neq 0$ and $\alpha_i = \mathbf{0}$ otherwise.

TABLE I
DATA DESCRIPTIONS.

Data set	Number of data	Dimensionality
Yale	165	3456
AT&T	400	2576
Umist	575	2576
AR	840	3072
XM2VTS	1180	4096
Coil20	1440	4096

Fig. 1
SOME SAMPLES FROM SIX BENCHMARK DATA SETS. THE IMAGES FROM THE FIRST ROW TO THE SIXTH ROW ARE: YALE, AT&T, UMIST, AR, XM2VTS, COIL20.

We use the following reconstruction error to measure the quality of dimensionality reduction methods:

$$e(m) = \frac{1}{n}\sum_{i=1}^{n}\left\|x_i^{org} - WW^T x_i\right\|_2, \tag{37}$$

where n is the number of training data, $W \in \Re^{d\times m}$ is the learned projection matrix by PCA, R1-PCA, PCA-L1, or PCA-L21, x_i^{org} and x_i are the i-th original non-occluded image and the i-th image used in the training respectively.

In the experiments, the projected dimension m varies from 21 to 69. The results of the reconstruction error by PCA, R1-PCA, PCA-L1 and PCA-L21 are shown in Figures 2-4. From the figures, we can see that PCA, R1-PCA and the proposed PCA-L21 are suitable for principal component analysis from the view of data reconstruction. When there are occlusions in the data, R1-PCA and PCA-L21 outperform PCA in terms of reconstruction error. We can also observed that in this experimental setting, PCA-L1 doesn't perform as good as PCA, which indicates that PCA-L1 is not closely connected to the minimization of reconstruction error, thus is not a good option for principal component analysis in some cases.

B. Reconstruction errors with noise images

In this experiment, two image data sets XM2VTS and Coil20 are used. For each data set, we add 10, 20, 30 percent images from the Palm image data set as the noise images, respectively. Some samples from the Palm data set are shown in Figure 5.

We use the following reconstruction error to measure the quality of dimensionality reduction methods:

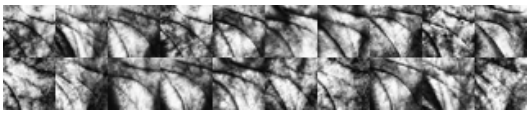

Fig. 5
TWENTY IMAGE SAMPLES FROM THE PALM IMAGE DATA SET.

$$e(m) = \frac{1}{n}\sum_{i=1}^{n}\left\|x_i^{org} - WW^T x_i^{org}\right\|_2, \tag{38}$$

where n is the number of training data (not including the noise image data from the Palm data set), $W \in \Re^{d\times m}$ is the learned projection matrix by PCA, R1-PCA, PCA-L1 or PCA-L21, x_i^{org} is the i-th original training data (not including the noise image data from the Palm data set). Under this experimental setting, if a method is robust to the data outliers, its reconstruction error should be smaller than other non-robust methods.

In the experiments, the projected dimension m varies from 21 to 69. The results of the reconstruction error by PCA, R1-PCA, PCA-L1 and PCA-L21 are shown in Figure 6.

We can see from the figures that, in this experimental setting, R1-PCA, PCA-L1 and the proposed PCA-L21 all outperform PCA in terms of reconstruction error, and our PCA-L21 consistently outperforms R1PCA, PCA-L1 and performs best in this case. The experimental results clearly indicates that the proposed PCA-L21 is more suitable for principal component analysis than the traditional PCA when there are outliers in the data.

VI. CONCLUSIONS

A principal component analysis with L21-norm maximization was proposed in this paper. The L21-norm maximization based PCA is theoretically connected to the minimization of the reconstruction error, and thus is more suitable for principal component analysis than the L1-norm maximization based PCA proposed in [11]. To avoid the greedy strategy used in [11] for solving the L1-norm maximization problem, we propose an efficient optimization algorithm to solve a more general L21-norm maximization problem, which is non-greedy and is guaranteed to converge to a local solution. Moreover, we extend our algorithm to solve the more general maximization problem which can derive solutions for many related statistical learning models. Experimental results on real world data sets show that the proposed method is effective for principal component analysis, and always obtains smaller reconstruction error than the related methods under the same reduced dimension.

ACKNOWLEDGMENT

This research was partially supported by nsf-iis 1117965, nsf-ccf 0830780, nsf-dms 0915228, nsf-ccf 0917274.

REFERENCES

[1] I. T. Jolliffe, *Principal Component Analysis,2nd Edition.* New York: Springer-Verlag, 2002.
[2] R. O. Duda., P. E. Hart., and D. G. Stork, *Pattern Classification.* Hoboken, NJ: Wiley-Interscience, 2000.
[3] A. Baccini, P. Besse, and A. de Faguerolles., "A L1-norm PCA and heuristic approach," in *Proceedings of the International Conference on Ordinal and Symbolic Data Analysis*, vol. 1, 1996, pp. 359–368.

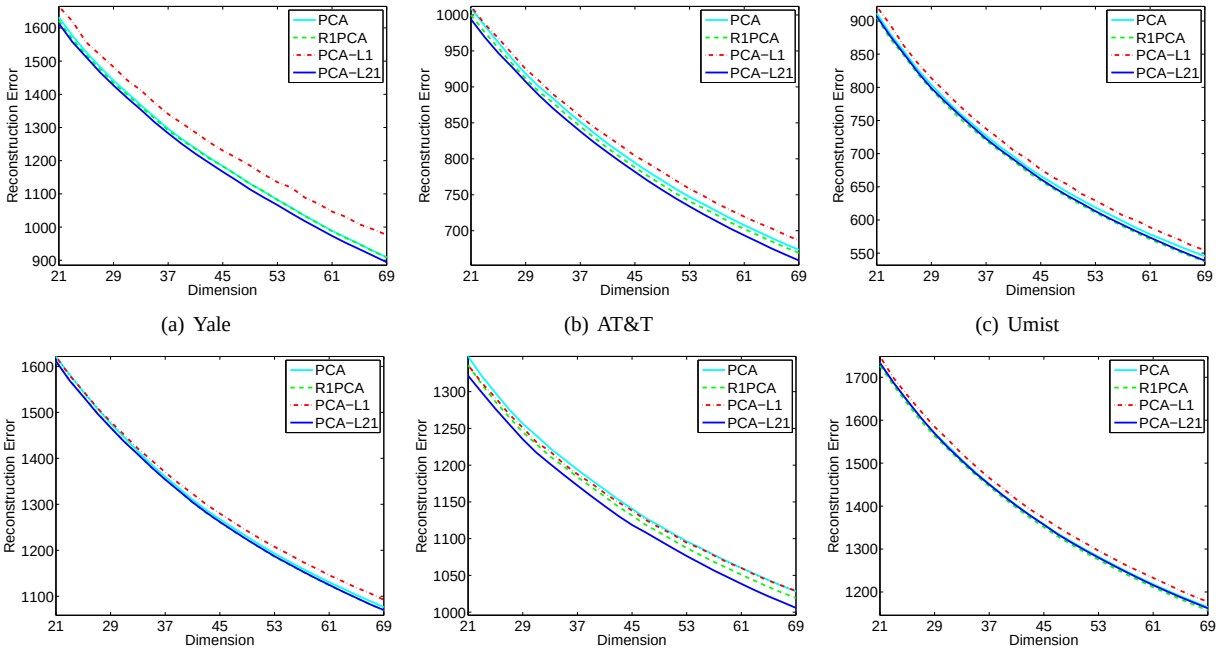

(a) Yale (b) AT&T (c) Umist

(d) AR (e) XM2VTS (f) Coil20

Fig. 2

RECONSTRUCTION ERRORS (CALCULATED BY EQ. (37)) UNDER DIFFERENT DIMENSIONS OBTAINED BY PCA, R1PCA, PCA-L1 AND PCA-L21, RESPECTIVELY. IN EACH DATA SET, 10 PERCENT IMAGES ARE RANDOMLY OCCLUDED WITH A RANDOMLY LOCATED SQUARE.

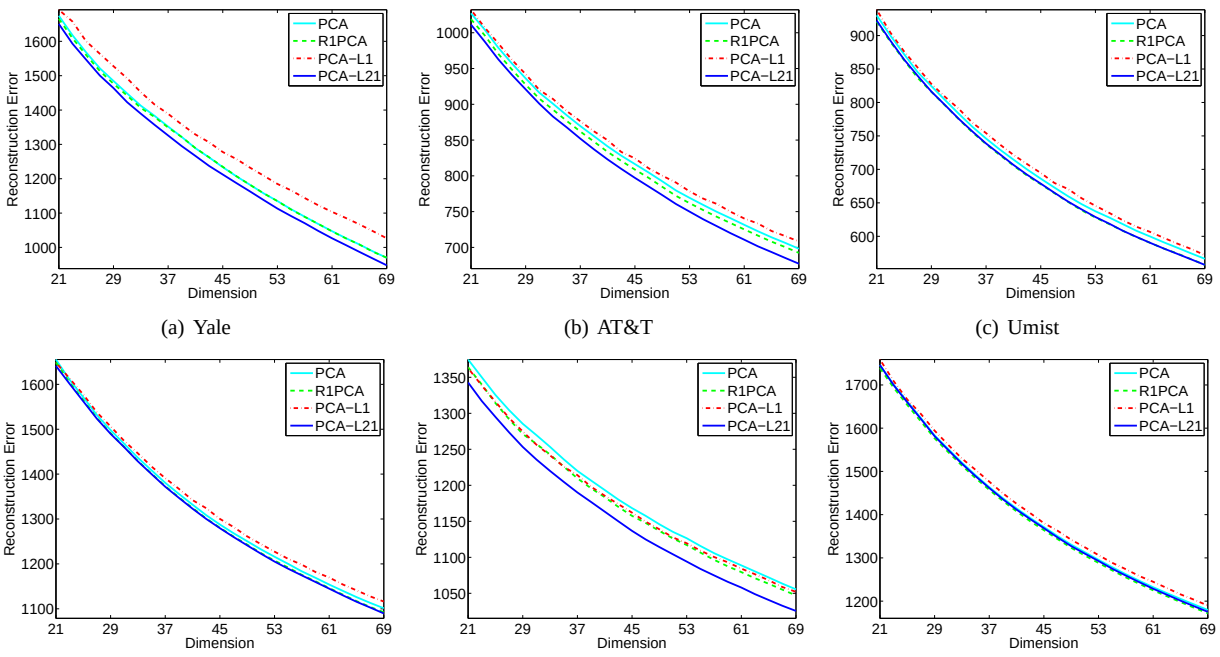

(a) Yale (b) AT&T (c) Umist

(d) AR (e) XM2VTS (f) Coil20

Fig. 3

RECONSTRUCTION ERRORS (CALCULATED BY EQ. (37)) UNDER DIFFERENT DIMENSIONS OBTAINED BY PCA, R1PCA, PCA-L1 AND PCA-L21, RESPECTIVELY. IN EACH DATA SET, 20 PERCENT IMAGES ARE RANDOMLY OCCLUDED WITH A RANDOMLY LOCATED SQUARE.

[4] H. Aanas, R. Fisker, K. Astrom, and J. Carstensen, "Robust factorization," *IEEE Transactions on PAMI*, vol. 24, no. 9, pp. 1215–1225, 2002.

[5] F. De La Torre and M. Black, "A framework for robust subspace learning," *International Journal of Computer Vision*, vol. 54, no. 1, pp. 117–142, 2003.

[6] Q. Ke and T. Kanade, "Robust L1 norm factorization in the presence of outliers and missing data by alternative convex programming," in *CVPR*, vol. 1, 2005, pp. 739–746.

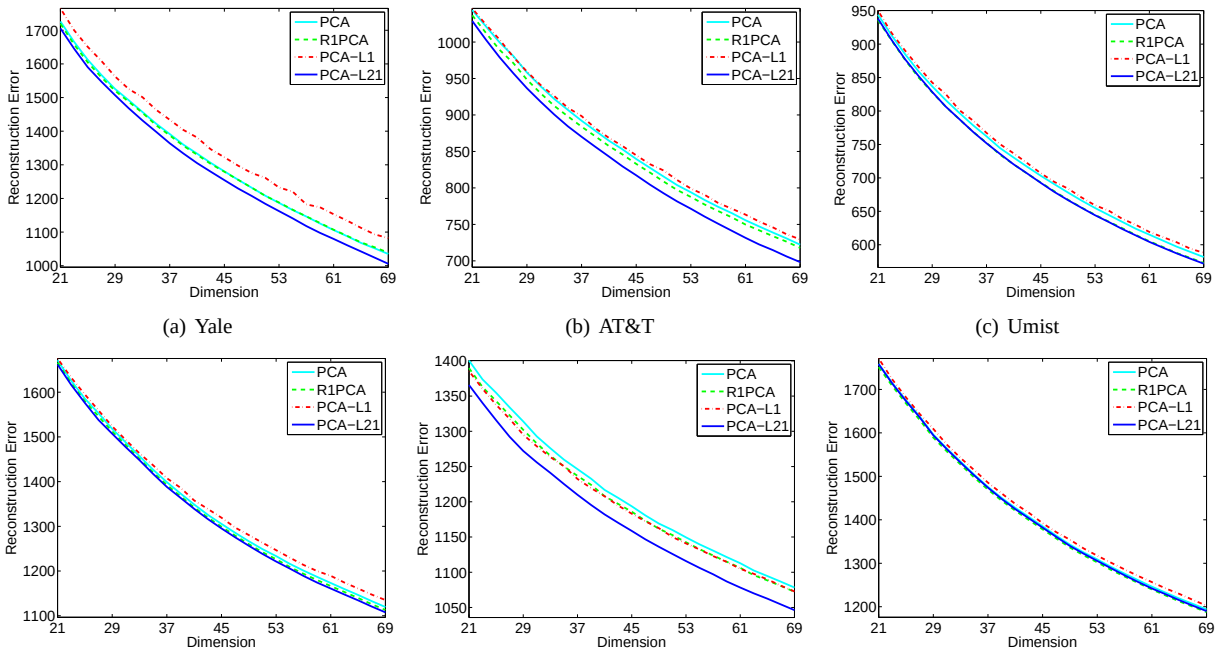

(a) Yale (b) AT&T (c) Umist

(d) AR (e) XM2VTS (f) Coil20

Fig. 4

RECONSTRUCTION ERRORS (CALCULATED BY EQ. (37)) UNDER DIFFERENT DIMENSIONS OBTAINED BY PCA, R1PCA, PCA-L1 AND PCA-L21, RESPECTIVELY. IN EACH DATA SET, 30 PERCENT IMAGES ARE RANDOMLY OCCLUDED WITH A RANDOMLY LOCATED SQUARE.

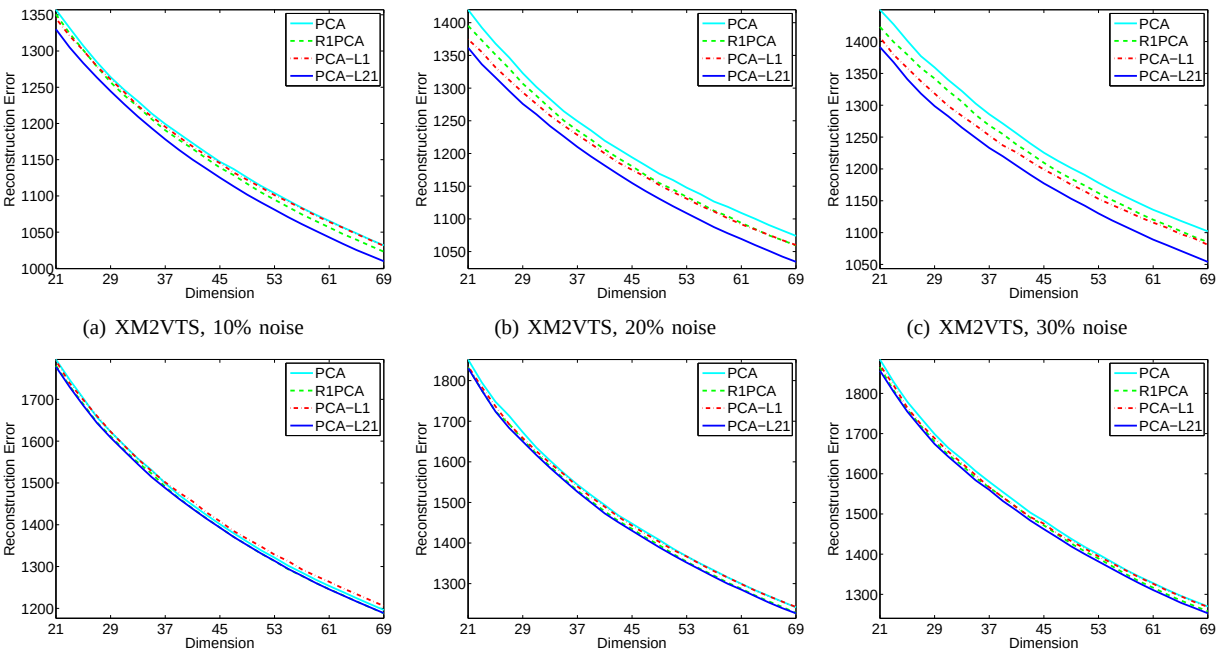

(a) XM2VTS, 10% noise (b) XM2VTS, 20% noise (c) XM2VTS, 30% noise

(d) Coil20, 10% noise (e) Coil20, 20% noise (f) Coil20, 30% noise

Fig. 6

RECONSTRUCTION ERRORS (CALCULATED BY EQ. (38)) UNDER DIFFERENT DIMENSIONS OBTAINED BY PCA, R1PCA, PCA-L1 AND PCA-L21, RESPECTIVELY. FIRST ROW INCLUDES THE RESULTS ON THE XM2VTS DATA SET WITH ADDING 10, 20 OR 30 PERCENT IMAGES FROM THE PALM DATA SET AS NOISE IMAGES. SECOND ROW SHOWS THE RESULTS ON THE COIL20 DATA SET WITH ADDING 10, 20 OR 30 PERCENT IMAGES FROM THE PALM DATA SET AS NOISE IMAGES.

[7] C. H. Q. Ding, D. Zhou, X. He, and H. Zha, "R1-PCA: rotational invariant L1-norm principal component analysis for robust subspace factorization," in *ICML*, 2006, pp. 281–288.

[8] J. Wright, A. Ganesh, S. Rao, and Y. Ma, "Robust principal component

analysis: Exact recovery of corrupted low-rank matrices via convex optimization," *NIPS*, 2009.
[9] G. H. Golub and C. F. van Loan, *Matrix Computations, 3rd Edition*. Baltimore, MD, USA: The Johns Hopkins University Press, 1996.
[10] H. Huang and C. H. Q. Ding, "Robust tensor factorization using r1 norm," in *CVPR*, 2008.
[11] N. Kwak, "Principal component analysis based on L1-norm maximization," *IEEE Transactions on PAMI*, vol. 30, no. 9, pp. 1672–1680, 2008.
[12] J. Galpin and D. Hawkins, "Methods of L1 estimation of a covariance matrix," *Computational Statistics & Data Analysis*, vol. 5, no. 4, pp. 305–319, 1987.
[13] X. Li, Y. Pang, and Y. Yuan, "L1-norm-based 2DPCA," *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, vol. 38, no. 4, 2010.
[14] Y. Liu, Y. Liu, and K. C. C. Chan, "Multilinear maximum distance embedding via l1-norm optimization," in *AAAI*, 2010.
[15] Y. Pang, X. Li, and Y. Yuan, "Robust tensor analysis with L1-norm," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 20, no. 2, pp. 172–178, 2010.
[16] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004.
[17] C. Zhang, F. Nie, and S. Xiang, "A general kernelization framework for learning algorithms based on kernel PCA," *Neurocomputing*, vol. 73, no. 4-6, pp. 959–967, 2010.
[18] B. Schölkopf, A. J. Smola, and K.-R. Müller, "Nonlinear component analysis as a kernel eigenvalue problem," *Neural Computation*, vol. 10, no. 5, pp. 1299–1319, 1998.
[19] L. D. Lathauwer, "Signal processing based on multilinear algebra," Ph.D. dissertation, Faculteit der Toegepaste Wetenschappen. Katholieke Universiteit Leuven, 1997.
[20] X. Yuan and T. Zhang, "Truncated power method for sparse eigenvalue problems," *Technical Report*, 2011.